Data processing apparatus with address redirection in response to periodic address patterns

The invention relates to a data processing apparatus and to a method of data processing.

From US patent No. 4,956,768 a data processing apparatus is known that provides for double buffering of data transferred between a processor and a plurality of outlets. Each outlet is provided with a pair of memories. The processor alternately writes to a first one of the memories and to a second one of the memories. When the processor writes to one memory the other memory is coupled to the outlet. Thus, writing from the processor and output to the outlet can proceed in parallel.

A processor associated with the outlet controls which of the memories is connected to the processor and which is connected to the outlet. US patent No. 4,956,768 does not describe how locations within the memories are addressed and under what conditions the role of the memories is switched.

Conventionally double buffering is used to provide decoupling between devices that produce and consume data from a data stream, respectively. A writing device alternately addresses one memory and another to write blocks of data. The reading device reads the block by addressing the memory that is not being addressed for writing. Usually, moreover, some form of signaling is required between the devices to indicate when the writing device switches from one block to another.

Amongst others, it is an object of the invention to provide for a form of double buffering communication between different data processing units in which double buffering is supported transparently for the data processing units.

The invention provides for a data processing apparatus according to Claim 1. According to the invention an independent switching unit controls which memory unit is connected to which data processing unit. Addresses from the data processing units are used to address locations in a memory unit selected by the switching unit, so that a given address

may address a location in different ones of the memory units at different times during execution of the same program, depending on the selection by the switching unit. The independent switching unit monitors the addresses supplied by at least one of the data processing units to detect repetitions in the pattern of addresses supplied by the processing unit. Upon detection of a repetition the switching unit switches the selection of the memory unit that is connected to the data processing unit.

Preferably, the criterion for detecting the repetitions is programmable, using for example detection of a repetition of addresses in a programmable range, or a programmable number of repetition, or a programmable combination of repetition of addresses from different ones of the processing units (e.g. alternating after detection of repetitions from both processing units, or alternating for each particular processing unit when a repetition is detected in the address pattern of that particular processing unit, optionally conditional on detection of a repetition by another processing unit after a preceding alternation of the address mapping from the particular one of the processing units.)

Various methods of detecting repetitions may be used, such as after detection of a repetition of an address received from a processing unit, or detection of a certain number of access operations within a certain range, or detection after use of all addresses in a certain range.

In general, the data processing apparatus will contain further memory units whose connections are not switched by the switching unit. Thus, only a subset of the range of addresses that a data processing unit may use to address memory addresses locations in the memory units that are connected via the switching unit. While the addresses in that subset are mapped alternately to different memory units, the remaining addresses are generally mapped to the same memory units. The switching unit preferably only monitors for repetition of addresses in the subset of addresses that address locations in the memory units that are connected to the data processing units via the switching unit. Thus, the alternations between different memory units are not directly dependent on patterns of addressing outside the memory units that are connected via the switching unit.

At least two data processing units and at least two memory units may be connected via the switching unit. However, the invention is easily scalable. Without deviating from the scope of the invention a greater number than two data processing unit and/or memory units may be connected, so that an address from a data processing unit can be mapped to any one of three or more memory units. In this case the switching unit may alternately connect three or more memory units to a data processing unit in a round-robin

fashion. Alternatively, the switching unit may be programmable so as to select which subset of the memory units is connected alternately to a specific data processing unit. Thus the data processing apparatus can be configured to provide flexible multiple buffered communication of more than one stream of data between more than two data processing units.

These and other objects and advantageous aspects of the data processing apparatus according to the invention will be described with reference to the following figures:

Fig. 1 shows a data processing apparatus,

Fig. 2 shows an embodiment of a switching control unit,

Fig. 3 shows a further embodiment of a switching control unit,

Fig. 3a shows another embodiment of a switching control unit, and

Fig. 4 shows a switching unit

Fig. 1 shows a data processing apparatus. The apparatus comprises processing units 10a,b, a plurality of memory units 18a-c, a switching unit 17 and a switching control unit 16. Each data processing unit 10a,b has connections to a respective address/control bus 14, 15 and a respective data bus 12, 13. The address bus 12 and the data bus 14 of a first processing unit 10a are shown coupled to a first port of the switching unit 17 and a first one of the memory units 18a. The address/control bus 13 and the data bus 15 of a second processing unit 10a are shown coupled to a second port of the switching unit 17. Furthermore, the address/control busses 14, 15 of the first and the second processing unit 10a,b are coupled to the switching control unit 16. The switching control unit 16 has a control output coupled to the switching unit 17. The switching unit 17 has third and fourth ports with connections for address/control and data bus lines to a second and a third one of the memory units 18a-c, respectively.

In operation the processing units 10a,b execute programs that include instructions for reading and/or writing data from and to memory locations. The instructions define the addresses of the relevant memory locations. In response to the instructions, the processing units 10a,b supply these addresses to the memory units 18a-c via the address/control busses 14, 15. Dependent on whether the instructions are read or write instructions, the processing units 10a,b also read data via the data busses 12, 13, or write data

via the data busses 12, 13, respectively. The memory units 18a-c that contain the location addressed by the addresses return data from the addressed locations to the data busses 12, 13 or store data from these data busses 12, 13 at the addressed locations.

A first and a second one of the memory units 18b,c contain locations that are addressed by the same addresses. Dependent on a control signal from the switching control unit 16, the switching unit 17 passes these addresses selectively either to the first or the second memory unit 18b,c. Similarly, the data corresponding to these addresses is passed to the selected memory unit 18b,c. Thus, dependent on a state of the switching control unit 16, the address from a processing unit 10a,b either addresses a location in the first memory unit 18b or in the second memory unit 18c. When the address from the first processing unit 10a is outside the range of addresses that address locations in the first and second memory units 18b,c, the address may address a third memory unit 18a directly, that is, not via the switching unit 17.

Although only a single directly addressed third memory unit 18a has been shown, it will be understood that a plurality of such directly connected memory units may in fact be present, some coupled to the address/control bus 14 and the data bus 12 of the first processing unit 10a, and others coupled to the address/control bus 15 and the data bus 13 of the second processing unit 10b.

The switching control unit 16 contains a state holding circuit (not shown), such as a status register, that retains state information which determines which of the memory units 18b,c is connected to the address/control bus 14, 15 and the data bus 12, 13 of which of the processing units 10a,b. The switching control unit 16 updates this state information in dependence on addresses received from the processing units 10a,b via the address/control busses 14, 15. The switching control unit 16 uses these addresses to detect the start of different periods of a periodic pattern in the addresses. Each time the switching control unit 16 detects the start of a period it updates the state information so that the addresses will subsequently be applied to a different memory unit 18b,c. Various ways of detecting the start of a period may be used.

Fig. 2 shows a first embodiment of the switching control unit 16 in its simplest form, wherein only one signal is generated for controlling the switching unit 17. In this embodiment the switching control unit 16 contains an address comparator 20a coupled to the address/control bus 14. The comparator 20a has an output coupled to a status register 22a which in turn has an output coupled to a control input of the switching unit 17 (not shown). In operation the comparator 20a compares addresses from the address/control bus 14 with a

set address. When the set address is detected, the comparator 20a causes the content of the status register 22a to toggle, which in turn causes the switching unit 17 to swap the memory units 18b,c that are coupled to the first and the second processing unit 10a,b respectively.

Another embodiment of the switching control unit 16 has the same structure as shown in Fig. 2, but in this embodiment the comparator 20a is a comparator that signals when the address is anywhere in a range of addresses that address locations in the first or the second memory unit 18b,c (e.g. by making two comparisons, testing for an address lower than an upper boundary address and higher than a lower boundary address, or by using only a more significant part of the address). In this embodiment the unit 22a is a counter that counts the number of times that addresses in the range are addressed, resets and updates a state register that controls the connections made by the switching unit 17 at the start of a new period when a certain number of addresses has been counted. Such a certain number may be a predetermined number, or a programmable number that is set by a program executed by one of the processing units 18b,c. In this embodiment a circuit that generates a memory unit enable signal for the memory units 18b,c (or chip enable if each memory unit 18b,c is made up of a memory chip) may be used as the comparator 20a, which in this case may be used to provide memory unit enable signals to either memory unit 18b,c, depending on the memory unit that has been selected.

Fig. 3 shows a further embodiment of the switching control unit 16. In this embodiment the switching control unit 16 contains a read modify write memory 30, a detector 32 and a toggle flip-flop 34. The address/control bus 14 is coupled to an address input of the read modify write memory and a data output of the read modify write memory 30 is coupled to an input of the detector 32. An output of the detector 32 is coupled to an input of the toggle flip-flop 34, which in turn has an output coupled to the output of the switching control unit 16. This output of the detector 32 is furthermore coupled to a reset input of the read modify write memory 30. The read modify write memory 30 has a respective location for each address value that can be used to address locations in the first and the second memory 18b,c.

In operation this embodiment of the switching control unit 16 detects the start of a new period of addressing by checking for repeated addressing of any address in the first or the second memory unit 18a,b. The read modify write memory 30 keeps information for each address value, indicating whether the address value has been used in a current period. The addresses that address locations in the first or the second memory 18b,c address locations in the read modify write memory 30 as well. Initially, at the start of a period the

detector 32 resets the content of the read modify write memory 30. Each time when such an address is received the content of the corresponding location in the read modify write memory 30 is set. The data that was previously stored at that location in the read modify write memory 30 is tested by the detector. If this data has been set, the detector 32 signals a repetition of the period, which causes the data content in the toggle flip-flop 34 to toggle and causes the content of the read modify write memory 30 to be reset. The data content of the toggle flip-flop 34 controls the connection made by the switching unit 17. This embodiment makes it possible to provide a more refined form of period detection, ignoring, for example, when certain addresses are not used in a certain period.

In a further embodiment the detector 32 may be replaced by a counter that counts how many times data read from the read modify write memory 30 indicated that an address is used for the first time in a period. In this embodiment this counter signals a new period (causing switching by the switching unit 17, a reset of the read modify write memory 30 and a new period of the counting process) when a certain count is exceeded. Thus, the switching control unit signals a new period when enough different locations (more than the certain number) have been addressed. The certain number may be predetermined or programmable by the processing units 10a,b, e.g. by means of a register coupled to at least one of the processing units for setting an initial value of the counter. This embodiment makes it possible to realize a more refined period detection, e.g. ignoring repeated access operations with the same address.

Although the embodiments described with reference to the Figs. 2 and 3 so far show detection of the period from addresses on one of the address/control busses 14 it should be understood that any of the address/control busses 14, 15 could be used, or that addresses from a combination of the busses could be used, switching for example, when the start of a new period has been detected in any one or in each of the processing units 10a,b. When the switching unit 17 permits simultaneous connection of more than one processing unit 10a,b to the same memory unit (e.g. in a time-interleaved manner or on a request arbitration basis), the switching control unit 16 may switch the bus connections of the processing units 10a,b independently of one another between the two memory units 18b,c. For this purpose the switching control unit 16 may be provided with two detection circuits, each for example of one of the types discussed in the context of the Figs. 2 or 3, each coupled to a respective data and address bus and each controlling to which memory unit 18b,c that address data bus is connected.

Fig. 3a shows an example of an embodiment of the switching control unit 16 in which addresses from a combination of busses are used. The embodiment contains a first and a second repetition detection unit 300a,b, a status comparator 302 and a control register 304. Each of the repetition detection units 300a,b may contain an address comparator, or a read-modify write memory with a detector as shown in the Figs. 2 and 3. Outputs of the repetition detection units 300a,b are coupled to the status comparator 302, which in turn has an update output coupled to the control register 304. The control register 304 has a control output coupled to the switching unit 17 (not shown).

In operation the repetition detection units 300a,b detect repetitions in the pattern of addressing from respective ones of the processing units, for example, as described in the context of the Figs. 2 and 3. On detection of a repetition the repetition detection units 300a,b send a signal to the status comparator 302. In a simple embodiment the status comparator 302 keeps status bits which are set when the respective repetitions are set and the status comparator toggles the control register 304 once both status bits have been set, clearing the status bits. In this embodiment access operations from a particular processing unit are preferably suspended once the status bit for that processing unit has been set.

In another embodiment, the status comparator 302 does not delay the update of the control register 304 until status bits have been set for all processing units, but generates, for example, an update to alternate mapping of addresses from each particular processing unit once a repetition occurs in the address pattern from that particular processing unit, unless no alternation of mapping of addresses from another processing unit has occurred since the last alternation of the mapping of the processing unit, in question.

Although only the use of addresses to control switching has been illustrated in order to provide a simple design, it should be understood that in more complex designs other control signals from the address/control bus 14, 15 may be used as well. For example, the switching control unit 16 may be designed so that it uses addresses only when used for reading or only when used for writing. This makes it possible to realize a more refined period detection, e.g. ignoring read operations.

Furthermore, although the simple switching control units shown in the Figs. 2 and 3 have been assumed to detect repetitions in a fixed range of addresses or of a fixed address, as an alternative a programmable range or a programmable address may be used, controlled, for example by using one or more programmable registers in the comparator 20a for defining a detection range, the register being coupled to at least one of the processing units, so that this processing unit can write values in these registers under program control.

Similarly, an address range detector may be added to the detector of Fig. 3, so as to detect for each addresses whether it is in a programmed range. When a repetition of an address is detected by means of the read-modify write memory, a repetition is signaled only if the address is in the programmed range. Thus, the criteria for the detection of repetitions of address patterns can be adjusted in dependence on the program used under program control. Similarly, the criteria for detection repetitions from a combination of processing units 10a,b may be program controlled.

Preferably, all the addresses used for the detection of repetitions in the embodiments described with reference to the Figs. 2 and 3 are only addresses in the range of addresses that address memory locations in the memory units 18b,c. However, in a further embodiment addresses outside that range may be used. Thus, it is possible to permit a large variation of access patterns with addresses in the memory units 18b,c that are connected via switching unit 17, while still permitting detection of repetition using the addresses outside the range of addresses of those memory units 18b,c.

Fig. 4 shows an embodiment of the switching unit 17. The embodiment contains selectable address/data bus drivers 40a,b, 42a,b. A first address/data bus driver 40a is connected between the address/control bus 14 and the data bus 12 of the first processing unit 10a (not shown), on the one side and to data/address connections for the first memory unit 18a (not shown) on the other side. A second address/data bus driver 40b is connected between the address/control bus 14 and the data bus 12 of the first processing unit 10a (not shown) on the one side and to data/address connections for the second memory unit 18b (not shown) on the other side. A third address/data bus driver 42a is connected between the address/control bus 15 and the data bus 13 of the second processing unit 10b (not shown) on the one side and to data/address connections for the first memory unit 18a (not shown), on the other side. A fourth address/data bus driver 42b is connected between the address/control bus 15 and the data bus 13 of the second processing unit 10b (not shown), on the one side and to data/address connections for the second memory unit 18b (not shown), on the other side.

The switching unit 17 has an input 44 for a control signal which is coupled to enable inputs of the address/data bus drivers 40a,b, 42a,b, so that either first and fourth address/data bus drivers 40a, 42b are enabled simultaneously or second and third address/data bus drivers 40b, 42a are enabled simultaneously, depending on a control signal from the switch control unit 16 (not shown). When enabled, the address/data bus drivers 40a,b, 42a,b pass data signals and address signals.

Without deviating from the scope of the invention more complex control signals may be used, for example, control signals that allow the first and second address/data bus drivers 40a,b to be controlled independently from the third and fourth address/data bus drivers 42a,b, provided that some sharing mechanism is provided that permits shared access to the memory units 18a,b (for example a time slot multiplexing mechanism, a priority mechanism or an arbitration mechanism).

The switching unit 17 can easily be expanded to support a greater number of processing units and or memory units. More processing units are supported by adding more address/data bus drivers; more memory units are supported by connecting more address/data bus drivers together. More than two memory units may be used, for example, to map addresses from the processing units 10a,b to different memories in a round-robin fashion. For this purpose, instead of the toggle flip-flops shown in the Figs. 2 and 3 cycling counters may be used for each of the processing units 10a,b so as to select which of the memory units 18a,b are addressed by each of the processing units 10a,b. This makes it more readily possible to switch the mapping of addresses from the individual processing units 10a,b once a repetition has been detected, without causing memory conflicts. Under program control different cycles of the memory units 18a-b may be selected, for example, alternating addressing of some addresses between a first pair of memory units and alternating addressing of other addresses between a second pair of memory units, or cycling mapping of the addresses through three of the memory units.

It will now be appreciated that the apparatus permits a processing unit 10a,b to address locations in respective ones of the memory units 18a,b with the same address. Control of the memory unit in which the location is addressed is exerted by a switching control unit that is external to the processing unit and that selects to switch in dependence on the detection of the start of a repetition of a periodic pattern. Although the same address from any one processing unit 10a,b addresses locations in different memory units 18a,b in dependence on the state of the switching control unit 16, it is not strictly necessary that locations in the memory units 18a,b are addressed by the same addresses from different ones of one processing unit 10a,b. Some address translation mechanism (if only suppression of a more significant part of the address that is not needed to distinguish addresses within the memory units 18a,b) may be included between the processing units 10a,b and the memory units 18a,b so that different addresses address the same locations, in dependence on the processing unit from which the address is supplied.